

Functional Choice and Non-significance Regions in Regression

Laurie Davies

Fakulät Mathematik

Universität Duisburg-Essen

`laurie.davies@uni-due.de`

Abstract

Given data y and k covariates x the problem is to decide which covariates to include when approximating y by a linear function of the covariates. The decision is based on replacing subsets of the covariates by i.i.d. normal random variables and comparing the error with that obtained by retaining the subsets. If the two errors are not significantly different for a particular subset it is concluded that the covariates in this subset are no better than random noise and they are not included in the linear approximation to y .

1 Introduction

1.1 Notation

Consider n measurements $\mathbf{y}_n = (y_1, \dots, y_n)^t$ of a variable y and for each y_i concomitant measurements of k covariables $x_j, j = 1, \dots, k$, given by $\mathbf{x}_i = (x_{i1}, \dots, x_{ik})^t$ forming an $n \times k$ matrix \mathbf{x}_n with j th column \mathbf{x}_j . A subset of the covariates will be denoted by a row vector $e = (e_1, \dots, e_k)$ with $e_j \in \{0, 1\}$ whereby $e_j = 1$ means that the j th covariate is included. A model e will be encoded as $\sum_{j=1}^k e_j 2^{j-1}$. The subset consisting of all covariates will be denoted by e_f . Given an e with $\sum_{j=1}^k e_j = k(e)$ the $n \times k(e)$ matrix with columns corresponding to those covariates with $e_j = 1$ will be denoted by $\mathbf{x}_n(e)$ with $\mathbf{x}_{i\cdot}(e)$, $\mathbf{x}_{\cdot j}(e)$ and $\mathbf{x}(e)$ having the corresponding interpretations. The empirical measure of the data will be denoted by \mathbb{P}_n

$$\mathbb{P}_n = \mathbb{P}_n((\mathbf{y}_n, \mathbf{x}_n)) = \frac{1}{n} \sum_{i=1}^n \delta_{y_i, \mathbf{x}_i}. \quad (1)$$

with the corresponding definition of $\mathbb{P}_{n,e}$ for any subset e . The L_1 and L_2 norms will be denoted by $\|\cdot\|_1$ and $\|\cdot\|_2$ respectively.

1.2 The problem

The problem to decide which if any of the covariates x_j influence the value of y . There are many proposals for doing this. Some such as AIC (Akaike (1973, 1974, 1981)) or BIC (Schwarz (1978)) require an explicit model such as

$$Y = \mathbf{x}^t \boldsymbol{\beta} + \varepsilon \quad (2)$$

where $\mathbf{x}^t = (x_1, \dots, x_k)$, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)^t$ and the errors ε are random variables with an explicit distribution. Others such as Lasso (Tibshirani (1996))

$$\operatorname{argmin}_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^n (y_i - \mathbf{x}_i \cdot \boldsymbol{\beta})^2 + \lambda \sum_{j=1}^k |\beta_j| \right\} \quad (3)$$

may or may not require an explicit model to determine the choice of the smoothing parameter λ .

The following is based on a simple idea. Let s^2 denotes the least sum of squares based on all covariates and for a given subset let S_e^2 denote the least sum of squares when all the covariates with $e_j = 0$ are replaced by i.i.d. $N(0, 1)$ random variables. If s^2 is not significantly small than S_e^2 the conclusion is that the omitted covariates are no better than random noise. To define ‘significantly’ the process is repeated a large number of times. For a given α , $\alpha = 0.95$ for example, s^2 is significantly smaller the S_e^2 if in at least $100\alpha\%$ the simulations $s^2 \leq S_e^2$. The P-value p_e is the proportion of simulations for which $S_e^2 < s^2$, so that the excluded covariates are significantly better than random noise if $p_e \leq 1 - \alpha$. A small values of p_e indicates that at least some of the omitted covariates are relevant. A large value of p_e indicates that in toto the omitted covariates are no better than random noise.

The method is not restricted to least squares regression. It can be equally well applied to L_1 regression or more generally to any measure of discrepancy $d(\mathbf{y}_n, \mathbf{x}_n)$.

As an example consider the stack loss data of Brownlee (1960). It is one of the data sets provided by R Core Team (2013) and is used in Koenker (2010). There are 21 observations with one dependent variable ‘Stack.Loss’ and the

three covariates ‘Air.Flow’, ‘Water.Temp’ and ‘Acid.Conc’ labelled from one to three. In the following the intercept will always be included. There are eight possible models. The results of an L_1 regression for the stack loss data are given in Table 1. The total computing time was 102 seconds using Koenker (2010). The only subset with a large P -value is the subset encoded as 3 which corresponds to $e = (1, 1, 0)$.

subset	0	1	2	3	4	5	6	7
P -value	0.000	0.015	0.000	0.231	0.000	0.007	0.000	1.000

Table 1: Encoded subsets and P -values for the stack loss data based on 5000 simulations.

1.3 Non-significance regions

Given a subset e of covariates the best linear approximation to the variable \mathbf{y}_n in the L_1 norm is

$$\mathbf{x}_n(e)\boldsymbol{\beta}_{1,n}(e) \tag{4}$$

where

$$\boldsymbol{\beta}_{1,n}(e) = \operatorname{argmin}_{\boldsymbol{\beta}(e)} \|\mathbf{y}_n - \mathbf{x}_n(e)\boldsymbol{\beta}(e)\|_1. \tag{5}$$

A single value is not sufficient to answer many questions of interest which require a range of plausible values. In frequentist statistics such a range is provided by a confidence region. This option is not available in the present context as a confidence region assumes that there is a ‘true’ value to be covered. The confidence region will be replaced by a non-significance region whose construction will be illustrated for the median.

Given data \mathbf{y}_n the median minimizes $s_1(\mathbf{y}_n) = \sum_{i=1}^n |y_i - \operatorname{med}(\mathbf{y}_n)|$. For any other value $m \neq \operatorname{med}(\mathbf{y}_n)$

$$\sum_{i=1}^n |y_i - \operatorname{med}(\mathbf{y}_n)| < \sum_{i=1}^n |y_i - m|$$

A value m will be considered as not being significantly different from the median

$\text{med}(\mathbf{y}_n)$ if the difference

$$\sum_{i=1}^n |y_i - m| - \sum_{i=1}^n |y_i - \text{med}(\mathbf{y}_n)|$$

is of the order attainable by a random perturbation of the y -values. More precisely if

$$\mathbf{P} \left(\inf_b \sum_{i=1}^n |y_i + bZ_i - m| < s_1(\mathbf{y}_n) \right) \geq 1 - \alpha. \quad (6)$$

The set of values m which satisfy (6) can be determined by simulations. For the \mathbf{y}_n of the stack loss data the 0.95-non-significance region is [11.86, 18.71] which can be compared with the 0.95-confidence region [11, 18] based on the order statistics. For any m the P-value $p(m)$ is defined as

$$p(m) = \mathbf{P} \left(\inf_b \sum_{i=1}^n |y_i + bZ_i - m| < s_1(\mathbf{y}_n) \right). \quad (7)$$

2 Choice of functional

The procedure expounded in the previous section makes no use of a model of the form (2). It solely based on the approximation of \mathbf{y}_n by a linear combination of the covariates as measured in the L_1 and L_2 norms. There is no mention of an error term. It therefore makes little sense to describe the procedure as one of model sense. It makes more sense to interpret it as one of functional choice. There does not seem to be any immediate connection with 'wrong model' approaches as in Berk et al. (2013) and Lindsay and Liu (2009).

For a given subset e the L_1 function $T_{1,e}$ is defined by

$$\begin{aligned} T_{1,e}(\mathbb{P}_n) &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \int |y - \mathbf{x}(e)^t \boldsymbol{\beta}(e)| d\mathbb{P}_n(y, \mathbf{x}(e)) \\ &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \|\mathbf{y}_n - \mathbf{x}_n(e) \boldsymbol{\beta}(e)\|_1 \end{aligned} \quad (8)$$

with the corresponding definition of the L_2 functional

$$\begin{aligned} T_{2,e}(\mathbb{P}_n) &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \int (y - \mathbf{x}(e)^t \boldsymbol{\beta}(e))^2 d\mathbb{P}_n(y, \mathbf{x}(e)) \\ &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \|\mathbf{y}_n - \mathbf{x}_n(e) \boldsymbol{\beta}(e)\|_2. \end{aligned} \quad (9)$$

More generally an M -functional $T_{\rho,e}$ can be defined as

$$\begin{aligned} T_{\rho,e}(\mathbb{P}_n) &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \int \rho \left(\frac{y - \mathbf{x}(e)^t \boldsymbol{\beta}(e)}{\sigma_n} \right) dP(y, \mathbf{x}(e)) \\ &= \operatorname{argmin}_{\boldsymbol{\beta}(e)} \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_i(e)^t \boldsymbol{\beta}(e)}{\sigma_n} \right). \end{aligned} \quad (10)$$

The function ρ is taken to be convex with a bounded first derivative. This is the case for the default choice in this paper namely the Huber ρ -function defined by

$$\rho_c(u) = \begin{cases} \frac{1}{2}u^2 & : |u| \leq c \\ c|u| - \frac{1}{2}c^2 & : |u| > c \end{cases} \quad (11)$$

where c is a tuning constant. The functional can be calculated using the iterative scheme described in Chapter 7.8.2 of Huber and Ronchetti (2009).

For reasons of equivariance (10) contains a scale parameter σ_n which may be external or part of the definition of T_ρ (see Chapter 7.8 of Huber and Ronchetti (2009)). The default choice in this paper is the Median Absolute Deviation of the residuals from an L_1 fit:

$$\sigma_n = \text{mad}(\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}_{1,n}(e_f)). \quad (12)$$

One use of M -functionals is to protect against outlying y -values. The choice (12) preserves this property.

2.1 L_1 regression

The best linear fit based on all covariates is determined by

$$T_{1,e_f}(\mathbb{P}_n) = \boldsymbol{\beta}_{1,n}(e_f) = \underset{\boldsymbol{\beta}}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_{i,\cdot}^t \boldsymbol{\beta}| = \underset{\boldsymbol{\beta}}{\text{argmin}} \|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}\|_1 \quad (13)$$

with mean sum of absolute deviations

$$s_{1,n}(e_f) = \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_{i,\cdot}^t \boldsymbol{\beta}_{1,n}(e_f)| = \|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}_{1,n}(e_f)\|_1. \quad (14)$$

Let \mathbf{Z}_n be a $n \times k$ matrix with elements Z_{ij} which are i.i.d. $N(0, 1)$. Given e replace the covariates with $e_j = 0$ by the Z_{ij} , that is, put $W_{i,j}(e) = x_{i,j}$ if $e_j = 1$ and $W_{i,j}(e) = Z_{ij}$ if $e_j = 0$. Denote the relevant matrices by $\mathbf{W}_n(e)$ and $\mathbf{Z}_n(e^c)$ and the empirical measure by $\tilde{\mathbb{P}}_{n,e}$. The best linear fit based on these covariates is determined

$$T_{1,e_f}(\tilde{\mathbb{P}}_{n,e}) = \tilde{\boldsymbol{\beta}}_{1,n}(e) = \underset{\boldsymbol{\beta}}{\text{argmin}} \|\mathbf{y}_n - \mathbf{W}_n(e) \boldsymbol{\beta}\|_1 \quad (15)$$

with mean sum of absolute deviations

$$S_{1,n}(e) = \|\mathbf{y}_n - \mathbf{W}_n(e) \tilde{\boldsymbol{\beta}}_{1,n}(e)\|_1 \quad (16)$$

The quantity $S_{1,n}(e)$ is a random variable. The P -value $p_n(e)$ is defined by

$$p_n(e) = \mathbf{P}(S_{1,n}(e) \leq s_{1,n}(e_f)). \quad (17)$$

There is no explicit expression for the P -values in the case of L_1 regression. They must be calculated using simulations as in Table 1. This results in four of the P -values being zero and so no comparison between them. A comparison can be obtained as follows. Simulate the distribution of

$$s_{1,n}(e) - S_{1,n}(e) \quad (18)$$

where

$$s_{1,n}(e) = \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^t \cdot \boldsymbol{\beta}_{1,n}(e)| = \|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}_{1,n}(e)\|_1 \quad (19)$$

and then approximate it by a Γ -distribution with the shape and scale parameters $sh(e)$ and $sc(e)$ estimated from the simulations as $\hat{sh}(e)$ and $\hat{sc}(e)$ respectively. The resulting estimated P -values are given by

$$\hat{p}_n(e) = 1 - \text{pgamma}(s_{1,n}(e) - s_{1,n}(e_f), \hat{sh}(e), \hat{sc}(e)). \quad (20)$$

The results for the stack loss data are given in Table 2 and may be compared with the P -values of Table 1.

functional	0	1	2	3	4	5	6	7
P -value	1.93e-7	1.41e-2	4.90e-4	2.32e-1	5.02e-9	7.43e-3	2.57e-4	1.00

Table 2: Encoded L_1 -functionals and P -values for the stack loss data based on 1000 simulations using the Γ -approximation (20),

Small P -values indicate that covariables have been omitted which have a significant effect on the dependent variable. This excludes the functionals encoded as 0, 1, 2, 4, 5, 6 although the functional encoded as 1 could possibly be retained. The functional 3 with P -value 0.232 omits the covariable Acid.Conc. As the functional 7 differs from 3 only through the inclusion of Acid.Conc the conclusion is that it contains a covariate which is little better than random noise. Thus an analysis of the P -values leads to the choice of the functional 3.

The interpretation of P -values and the choice of functional will be considered in greater detail in Sections 2.5 and 2.6 respectively.

The second running example is the low birth weight data of Hosmer and Lemeshow (1989) with $n = 189$ and $k = 9$. The dependent variable is the weight of the child at birth. The nine covariates range from the weight and age of the mother to hypertension and indicators of race. There are in all 512 different functionals. In the context of model choice it is considered in Claeskens and Hjort (2003).

For this data set the computing time using 1000 simulations is about 50 minutes. This can be reduced by a factor of about ten by approximating the modulus function $|x|$ by the Huber ρ -function (11) with a small value of the tuning constant c , for example $c = 0.01$ (see Section 2.2). Care must be taken in interpreting the decrease in computing time as the L_1 -functional was calculated using package Koenker (2010) whereas the program for the M -functional was written entirely in Fortran using the algorithm given in Chapter 7.8 of Huber and Ronchetti (2009) (see also Dutter (1977*b*) and Dutter (1977*a*)) and the pseudo-random number generator *ran2* (see Press et al. (2003)). A pure Fortran program for the L_1 -functional may be much faster (see Koenker and Portnoy (1997)).

2.2 M -regression functionals

The M -functionals can be treated in the same manner as the L_1 functional but with the added advantage that for large values of the tuning constant c in (11)

there exist asymptotic approximations for the P -values. On writing

$$T_{\rho, e_f}(\mathbb{P}_n) = \beta_\rho(e_f) = \operatorname{argmin}_{\beta} \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_{i \cdot}^t \beta}{\sigma_n} \right) \quad (21)$$

$$s_{\rho, n}(e_f) = \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_{i \cdot}^t \beta_\rho(e_f)}{\sigma_n} \right) \quad (22)$$

$$T_{\rho, e}(\mathbb{P}_n) = \beta_\rho(e) = \operatorname{argmin}_{\beta} \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_{i \cdot}(e)^t \beta}{\sigma_n} \right) \quad (23)$$

$$s_\rho(e) = \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_{i \cdot}(e)^t \beta_\rho(e)}{\sigma_n} \right) \quad (24)$$

$$T_{\rho, e}(\tilde{\mathbb{P}}_n) = \tilde{\beta}_\rho(e) = \operatorname{argmin}_{\beta(e)} \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{W}_{i \cdot}(e)^t \beta(e)}{\sigma_n} \right) \quad (25)$$

$$S_{\rho, n}(e) = \frac{1}{n} \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{W}_{i \cdot}(e)^t \tilde{\beta}_\rho(e)}{\sigma_n} \right) \quad (26)$$

a second order Taylor expansion gives

$$S_{\rho, n}(e) \approx s_{\rho, n}(e) - \frac{1}{2} \frac{\left(\frac{1}{n} \sum_{i=1}^n \rho^{(1)} \left(\frac{r_i(e)}{\sigma_n} \right)^2 \right) \chi_{k-k(e)}^2}{\frac{1}{n} \sum_{i=1}^n \rho^{(2)} \left(\frac{r_i(e)}{\sigma_n} \right)} \quad (27)$$

where $\rho^{(1)}$ and $\rho^{(2)}$ are first and second derivatives of ρ respectively, $r_i(e) = y_i - \mathbf{x}_{i \cdot}(e)^t \beta(e)$ and $\chi_{k-k(e)}^2$ is a chi-squared random variable with $k - k(e)$ degrees of freedom. The inequality $S_{\rho, n}(e) \leq s_{\rho, n}(e_f)$ is asymptotically equivalent to

$$\frac{\left(\frac{2}{n} \sum_{i=1}^n \rho^{(2)} \left(\frac{r_i(e)}{\sigma_n} \right) \right) (s_{\rho, n}(e) - s_{\rho, n}(e_f))}{\frac{1}{n} \sum_{i=1}^n \rho^{(1)} \left(\frac{r_i(e)}{\sigma_n} \right)^2} \leq \chi_{k-k(e)}^2 \quad (28)$$

with asymptotic P -value

$$p_n(e) \approx 1 - \text{pchisq} \left(\frac{\left(\frac{2}{n} \sum_{i=1}^n \rho^{(2)} \left(\frac{r_i(e)}{\sigma_n} \right) \right) (s_{\rho, n}(e) - s_{\rho, n}(e_f))}{\frac{1}{n} \sum_{i=1}^n \rho^{(1)} \left(\frac{r_i(e)}{\sigma_n} \right)^2}, k - k(e) \right). \quad (29)$$

As the tuning constant c tends to zero the terms $\rho_c^{(1)} \left(\frac{r_i(e)}{\sigma_n} \right)^2$ and $\rho_c^{(2)} \left(\frac{r_i(e)}{\sigma_n} \right)$ become one and zero respectively and the approximation breaks down.

The results for the stack loss data with $c = 1.5$ are given in Table 3 and may be compared with those given in Table 1 for the L_1 -functional.

functional	0	1	2	3	4	5	6	7
P -value	0.000	0.012	0.000	0.293	0.000	0.006	0.000	1.000
P -value	3.19e-7	1.23e-2	4.40e-4	3.03e-1	3.79e-8	5.74e-3	3.59e-5	1.00
P -value	1.89e-6	9.96e-3	1.81e-3	2.33e-1	4.63e-11	2.67e-3	6.66e-5	1.00

Table 3: Encoded M -functionals ($c=1.5$) and P -values for the stack loss data based on 5000 simulations: first row the raw values, second row the values based on the Γ -approximation (20), third row the values based on the asymptotic approximation (29).

2.3 Least squares regression

The L_2 -regression functionals are a special case of the M -functionals for a sufficiently large tuning constant c . The P -values can either be estimated directly using simulations or using the Γ -approximation (20) or using the asymptotic approximation (29) which takes the form

$$p_n(e) \approx 1 - \text{pchisq} \left(\frac{n(\|\mathbf{y}_n - \mathbf{x}_n(e)\beta_{2,n}(e)\|_2^2 - \|\mathbf{y}_n - \mathbf{x}_n\beta_{2,n}\|_2^2)}{\|\mathbf{y}_n - \mathbf{x}_n(e)\beta_{2,n}(e)\|_2^2}, k - k(e) \right). \quad (30)$$

If a lower bound α is given for $p_n(e)$ then (30) is asymptotically equivalent to the F -test in the linear regression model for testing the null hypothesis that the coefficients of the covariates not included are zero.

The P -values for the stack loss data are given in Table 4.

functional	0	1	2	3	4	5	6	7
P -value	0.000	0.011	0.000	0.427	0.000	0.005	0.000	1.000
P -value	1.89e-4	1.07e-2	1.28e-4	4.35e-1	2.80e-5	4.42e-3	4.24e-4	1.00
P -value	2.53e-4	1.51e-2	1.04e-4	3.11e-1	8.14e-5	4.69e-3	2.20e-4	1.00

Table 4: Encoded L_2 -functionals and P -values for the stack loss data based on 5000 simulations: first row the raw values, second row the values based on the Γ -approximation (20), third row the values based on the asymptotic approximation (30).

2.4 Non-linear regression

The ideas can be applied mutatis mutandis to non-linear regression

$$T_{nl,1,e}(\mathbb{P}_n) = \operatorname{argmin}_{\boldsymbol{\beta}(e)} \|\mathbf{y}_n - g(\mathbf{x}_n(e), \boldsymbol{\beta}(e))\|_1 \quad (31)$$

with corresponding definitions for $T_{nl,\rho,e}$ and $T_{nl,2,e}$. The computational cost is much higher so that only small values of k are possible.

2.5 Lower bounds for P -values

In contrast to AIC and BIC the P -values do not order the different functionals. One possibility is to choose a cut-off value $p_0(n, k)$ for $p_n(e)$ and consider only those functionals T_e with $p_n(e) \geq p_0(n, k)$. A possible value for $p_0(n, k)$ can be obtained by considering the size of the P -values when all covariates are noise, $\mathbf{x}_n = \tilde{\mathbf{Z}}_n$. For each such $\mathbf{x}_n = \tilde{\mathbf{Z}}_n$ the minimum value of $p_n(e)$ over all e can be calculated and then simulated for different $\tilde{\mathbf{Z}}_n$. The α quantile with for example $\alpha = 0.05$ can then be taken as the value of $p_0(n, k) = p_0(n, k, \alpha)$.

The minimum of the $p(e)$ -values can only be determined by simulation and then further simulations are required in order to determine the quantiles of the minimum values. For L_1 -functionals for the stack loss data with $n = 21$ and $k = 3$ the time required with 1000 simulations for each $p(e)$ and 2000 simulations for the quantile was 10 minutes using the approximation to $|x|$ based on the Huber ρ -function (11) with tuning constant $c = 0.01$. The results using the Γ -approximation (20) are given in the first line of Table 5.

The computational time for the low birth weight data with $k = 9$ was considerably higher. The time required for 1000 simulations for the minimum values of the $p_n(e)$ each of which was also based on 500 simulations was 34 hours. The results using the Γ -approximation (20) are given in the second line of Table 5.

The corresponding p_0 values for the M -functional and the L_2 -functional are based on 500 simulations with 250 simulations for each $p_n(e)$ value. The computing time for the low birth weight data was 2 1/2 hours in each case. The results are given in the lines 3-6 of Table 5.

For the M -functionals with a not too smaller a value of c in (11) and the L_2 -functionals use can be made of the approximations (29) and (30) respectively

L_1 -functional

$$\begin{array}{lll} p_0(21, 3, 0.01) = 0.00368 & p_0(21, 3, 0.05) = 0.0155 & p_0(21, 3, 0.10) = 0.0340 \\ p_0(189, 9, 0.01) = 0.00044 & p_0(189, 9, 0.05) = 0.0031 & p_0(189, 9, 0.10) = 0.0068 \end{array}$$

L_2 -functional

$$\begin{array}{lll} p_0(21, 3, 0.01) = 0.0055 & p_0(21, 3, 0.05) = 0.0193 & p_0(21, 3, 0.10) = 0.036 \\ p_0(189, 9, 0.01) = 0.00011 & p_0(189, 9, 0.05) = 0.0020 & p_0(189, 9, 0.10) = 0.0056 \end{array}$$

Table 5: The P -values $p_0(n, k, \alpha)$ for the L_1 - and L_2 -functionals.

which allow simulations for larger values of k . The computational load can be further reduced as follows. Without loss of generality suppose $\|\mathbf{y}_n\|_2 = 1$. As all the $\mathbf{x}_n = \tilde{\mathbf{Z}}_n$ are standard Gaussian white noise random variables it follows for the L_2 -functionals

$$\frac{\|\mathbf{Z}_n(e^c)^t(\mathbf{y}_n - \mathbf{x}_n(e)\beta_{2,n}(e))\|_2^2}{\|\mathbf{y}_n - \mathbf{x}_n(e)\beta_{2,n}(e)\|_2^2} \approx \|\mathbf{Z}_n(e^c)^t \mathbf{y}_n\|_2^2 \approx \sum_j \|\mathbf{Z}_n(e_j^c)^t \mathbf{y}_n\|_2^2 \quad (32)$$

where $e^c = \sum_j e_j^c$ and each e_j^c has only one element not equal to zero. Furthermore the $\mathbf{Z}_n(e_j^c)^t \mathbf{y}_n$ are independent $N(0, 1)$ random variables. This yields the asymptotic approximation

$$\tilde{p}(e) = 1 - \text{pchisq}\left(\sum_{j \in S} \chi_1^2(j), |S|\right) \quad (33)$$

where the $\chi_1^2(j), j = 1, \dots, k$ are independent χ^2 random variables with one degree of freedom. Taking the minimum over e , simulating sets of χ^2 random variables and then taking the α quantile gives a value $\tilde{p}_0(k, \alpha)$. It is only necessary to perform the simulations one for each value of k . The $\tilde{p}_0(k, \alpha)$ can be approximated by

$$\tilde{p}_0(k, \alpha) \approx \exp(c_1(k) + c_2(k) \log(\alpha) + c_3(k) \log(\alpha)^2) \quad (34)$$

for $\alpha < 0.5$ (see Chapter 2.9 of Davies (2014)) and can be used in place of the $p_0(n, k, \alpha)$. For $k = 17$ the time required on a standard laptop was 7 hours 42 minutes. As an example

$$\tilde{p}_0(9, 0.01) = 0.00028, \tilde{p}_0(9, 0.05) = 0.0025, \tilde{p}_0(9, 0.10) = 0.0059.$$

The results compare well with those based on simulations as given in Table 5 and their computing costs are essentially zero. This suggests that they can be used as guidelines when simulations are too expensive.

2.6 Choosing functionals

In view of the interpretation of the P -value of a functional the first step is to decide on a cut-off value p_0 and then restrict consideration to those functionals T_e with $p_n(e) > p_0$. A possible choice of p_0 is $p_0 = p_0(n, k, \alpha)$ and this will be done below. The choice may be further restricted by requiring that for each such e and for all e' with $e' < e$ pointwise all the omitted covariates are relevant with respect to e . More precisely $p_n(e', e) < p_0(n, k(e), \alpha)$ for all $e' < e$ where $p_n(e', e)$ is the P -value of e' calculated with respect to the covariates $\mathbf{x}(e)$. A final choice may be made by choosing that functional T_e with the highest $p_n(e)$ -value.

The results of applying the above strategy to the stack loss data with the L_1 -norm are as follows. Taking the cut-off value to be $p_0(21, 3, 0.01) = 0.00368$ (based on Table 5) the first step results in the functionals based on $e_1 = (1, 0, 0)$, $e_2 = (1, 1, 0)$ and $e_3 = (1, 0, 1)$. The second step eliminates the functionals based on e_2 and e_3 . The choices $\alpha = 0.05$ and $\alpha = 0.1$ both lead to $e = e_1$ with $p_n(e_1) = 0.232$.

The results for the low birth weight data are as follows. The first step with $\alpha = 0.01$ leads to 379 functionals. The second step results in the single functional encoded as 260 with just two variables and a P -value of $4.25e-3$. Putting $\alpha = 0.05$ results in 221 functionals after the first step. The second step yields the five functionals encoded as 292, 166, 260, 36 and 60. The functional 292 has the highest P -value with $p_n(e) = 0.074$. Finally the choice $\alpha = 0.1$ results in 149 functionals after the first step. The second step reduces this to the seven functionals 308, 292, 166, 262, 38, 52 and 260 of which the functional encoded as 308 has the highest P -value equal to 0.321.

The strategy described above ‘guarantees’ that all included covariates are relevant. If it is more important not to exclude covariates which may have an influence at the possible cost of including some irrelevant covariates then this may be done by increasing the value of α in $p_0(n, k, \alpha)$ or by simply specifying some cut-off level p_0 judged to be appropriate.

Although AIC and BIC list the models in order of preference they give no indication as to whether any of the models under consideration is an adequate approximation to the data or not. Presumably this is the responsibility of the user before applying the criterion. The first ten models for the birth weight data

based on BIC are encoded as

$$308, 310, 436, 438, 294, 292, 182, 316, 422, 309. \quad (35)$$

The functionals obtained using the P -values strategy are encoded as 260, 292 and 308. Their positions in the BIC list are 66, 6 and 1 respectively. The second model on the BIC list is encoded as 310 and includes the additional covariate ‘weight of mother’ compared to the 308 model. If one uses R Core Team (2013) to do an L_1 regression based on the covariates corresponding to 310 the 95% confidence interval for ‘weight of mother’ includes zero. This may be interpreted as a non-significant effect given the other covariates. This interpretation is consistent with the P -value strategy with $\alpha = 0.1$ where the encoded value 310 is not included in the list (35). The reason is that the P -value for the functional excluding ‘weight of mother’ has a P -value of 0.112 which exceeds $p_0(189, 5, 0.1) = 0.017$.

This shows that models may be high in the BIC list although they contain variables which are not significantly better than random noise. This can be made more explicit by replacing the all covariates by random noise and using simulations to determine how often a model containing a random noise covariate is first on the BIC list. This was simulated 500 times with the weight of the child as the dependent variable. This happened in 43% of the cases. With $\alpha = 0.1$ the P -value strategy is calibrated to do this in $100\alpha\% = 10\%$ of the cases. The simulations resulted in 9%.

3 Non-Significance regions

3.1 The median and M -functionals

The 0.95-non-significance region for the median of the stack loss data was defined and calculated in Section 1.3 with the result [11.86, 18.71]. In general the α -non-significance region is defined by \mathbf{y}_n is

$$\begin{aligned} \mathcal{NS}(\mathbf{y}_n, \text{med}, \alpha) \\ = \left\{ m : \sum_{i=1}^n |y_i - m| - \sum_{i=1}^n |y_i - \text{med}(\mathbf{y}_n)| \leq \text{ql1}(\alpha, m, \mathbf{y}_n) \right\} \end{aligned} \quad (36)$$

where $\text{ql1}(\alpha, m, \mathbf{y}_n)$ is the α -quantile of

$$\sum_{i=1}^n |y_i - m| - \inf_b \sum_{i=1}^n |y_i - m - bZ_i| \quad (37)$$

and the Z_i are standard Gaussian white noise.

The non-significance region (36) can be calculated as follows. Put

$$f(m, \alpha, \mathbf{y}_n) = \text{ql1}(\alpha, m, \mathbf{y}_n) - \sum_{i=1}^n |y_i - m| + \sum_{i=1}^n |y_i - \text{med}(\mathbf{y}_n)| \quad (38)$$

and note that $f(\text{med}(\mathbf{y}_n), \alpha, \mathbf{y}_n) \geq 0$. Now determine an order statistic $y_{(nl)}$ with $nl = \text{qbinom}((1 - \beta)/2, n, 0.5)$ for a suitably large β such that $f(y_{(nl)}, \alpha, \mathbf{y}_n) < 0$. Interval bisection combined with simulations can now be used to find an approximate solution m_{lb} of $f(m, \alpha, \mathbf{y}_n) = 0$. This gives a lower bound and the same process can be used to get an upper bound m_{ub} to give $\mathcal{NS}(\mathbf{y}_n, \text{med}, \alpha) = [m_{\text{lb}}, m_{\text{ub}}]$.

Non-significance regions for M -functionals T_ρ are defined analogously by replacing (36) by

$$\begin{aligned} & \mathcal{NS}(\mathbf{y}_n, T_\rho, \alpha) \\ &= \left\{ m : \sum_{i=1}^n \rho\left(\frac{y_i - m}{\sigma_n}\right) - \sum_{i=1}^n \rho\left(\frac{y_i - T_\rho(\mathbb{P}_n)}{\sigma_n}\right) \leq \text{qrho}(\alpha, m, \mathbf{y}_n) \right\} \end{aligned} \quad (39)$$

where $\text{qrho}(\alpha, m, \mathbf{y}_n)$ is the α -quantile of

$$\sum_{i=1}^n \rho\left(\frac{y_i - m}{\sigma_n}\right) - \inf_b \sum_{i=1}^n \rho\left(\frac{y_i - m - bZ_i}{\sigma_n}\right), \quad (40)$$

the Z_i are standard Gaussian white noise and σ_n is a scale functional whose default value in this situation is (12).

For smooth functions ρ an asymptotic expression for the non-significance region is available. Let $\rho^{(1)}$ and $\rho^{(2)}$ denote the first and second derivative of ρ . A Taylor series expansion results in

$$\mathcal{NS}(\mathbf{y}_n, T_\rho, \alpha) \approx \left\{ m : |T_\rho(\mathbb{P}_n) - m| \leq \text{qnorm}((1 + \alpha)/2) \sigma_n \sqrt{v(T_\rho, \mathbb{P}_n)/n} \right\} \quad (41)$$

where

$$v(T_\rho, \mathbb{P}_n) = \frac{\frac{1}{n} \sum_{i=1}^n \rho^{(1)}\left(\frac{y_i - T_\rho(\mathbb{P}_n)}{\sigma_n}\right)^2}{\left(\frac{1}{n} \sum_{i=1}^n \rho^{(2)}\left(\frac{y_i - T_\rho(\mathbb{P}_n)}{\sigma_n}\right)\right)^2}. \quad (42)$$

This latter expression is well known in robust statistics and corresponds to the asymptotic variance of an M -location functional: the non-significance region (41) is the corresponding α -confidence region for the ‘unknown’ $T_\rho(P)$. In the special case $\rho(u) = u^2/2$ (41) is the asymptotic α -confidence region for the mean based on Gaussian errors.

3.2 L_1 regression

The idea carries over to the L_1 regression functional. For any β put

$$\Gamma(\mathbf{y}_n, \mathbf{x}_n, \beta, \mathbf{Z}_n) = \|\mathbf{y}_n - \mathbf{x}_n\beta\|_1 - \inf_{\mathbf{b}} \|\mathbf{y}_n - \mathbf{x}_n\beta - \mathbf{Z}_n\mathbf{b}\|_1 \quad (43)$$

and denote the α -quantile of $\Gamma(\mathbf{y}_n, \mathbf{x}, \beta, \mathbf{Z}_n)$ by $\text{q1}(\alpha, \beta, \mathbf{y}_n, \mathbf{x}_n)$. An α -non-significance region is then defined as

$$\mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_1) = \{\beta : \|\mathbf{y}_n - \mathbf{x}_n\beta\|_1 - \|\mathbf{y}_n - \mathbf{x}_n\beta_{1,n}\|_1 \leq \text{q1}(\alpha, \beta, \mathbf{y}_n, \mathbf{x}_n)\} \quad (44)$$

where $\beta_{1,n} = T_1(\mathbb{P}_n)$.

As it stands the non-significance region is difficult to calculate as it requires a grid of values for the possible values of β and the values of $\text{q1}(\alpha, \beta, \mathbf{y}_n, \mathbf{x}_n)$ have to be estimated using simulations. If the quantiles are largely independent of the β -values then $\text{q1}(\alpha, \beta_1, \mathbf{y}_n, \mathbf{x}_n)$ can be used with a large reduction in computation. Section 3.5 contains some asymptotics which suggest that the independence may hold for large sample sizes n . The defining inequality in (44) will still have to be checked over a grid of values.

Most software packages provide only confidence regions for the individual components of β . Corresponding component wise non-significance regions can be defined with a large reduction in the computational overload. For the first component β_1 of $T_1(\mathbb{P}_n)$ the α -non-significance region is given by

$$\mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_{1,1}) = \left\{ \beta_1 : \inf_{\beta_2, \dots, \beta_k} \left\| \mathbf{y}_n - \mathbf{x}_{\cdot 1}\beta_1 - \sum_{j=2}^k \mathbf{x}_{\cdot j}\beta_j \right\|_1 - \left\| \mathbf{y}_n - \mathbf{x}_n\beta_{1,n} \right\|_1 \leq \text{q1}(\alpha, \beta_1, \mathbf{y}_n, \mathbf{x}_n) \right\} \quad (45)$$

where $q1(\alpha, \beta_1, \mathbf{y}_n, \mathbf{x}_n)$ is the α -quantile of

$$\begin{aligned} \Gamma_1(\mathbf{y}_n, \mathbf{x}_n, \beta_1, Z_{.1}) &= \inf_{\beta_2, \dots, \beta_k} \left\| \mathbf{y}_n - \mathbf{x}_{.1}\beta_1 - \sum_{j=2}^k \mathbf{x}_{.j}\beta_j \right\|_1 - \\ &\quad \inf_{b_1, \beta_2, \dots, \beta_k} \left\| \mathbf{y}_n - \mathbf{x}_{.1}\beta_1 - \sum_{j=2}^k \mathbf{x}_{.j}\beta_j - \mathbf{Z}_{.1}b_1 \right\|_1 \end{aligned} \quad (46)$$

The non-significance intervals of the stack loss data and for comparison the 0.95-confidence intervals are given in Table 6.

	Air.Flow	Water.Temp	Acid.Conc
Non-sig. intervals (45)	(0.552,1.082)	(0.225,1.603)	(-0.345,0.102)
L_1 confidence intervals	(0.509,1.168)	(0.272,3.037)	(-0.278,0.015)

Table 6: First line: 0.95-non-significance intervals for the stack loss data. Second line: 0.95-confidence intervals produced by Koenker (2010) for the default choice ‘se=rank’.

3.3 M -regression functionals

Non-significance regions for M -regression functionals are defined in the same manner as for L_1 regression. Just as in Section 2.2 the computational burden can be reduced for large n by using the asymptotic expressions. These result in

$$\begin{aligned} \mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_\rho) &= \left\{ \boldsymbol{\beta} : \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}}{\sigma_n} \right) - \sum_{i=1}^n \rho \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}_\rho}{\sigma_n} \right) \leq \right. \\ &\quad \left. \frac{\text{qchisq}(\alpha, k)}{2} \frac{\sum_{i=1}^n \rho^{(1)^2} \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}}{\sigma_n} \right)}{\sum_{i=1}^n \rho^{(2)} \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}_\rho}{\sigma_n} \right)} \right\} \end{aligned} \quad (47)$$

where $\boldsymbol{\beta}_\rho = T_\rho(\mathbb{P}_n)$. This can be further simplified to

$$\begin{aligned} \mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_\rho) &= \left\{ \boldsymbol{\beta} : (\boldsymbol{\beta} - \boldsymbol{\beta}_\rho)^t \mathbf{x}_n^t \mathbf{x}_n (\boldsymbol{\beta} - \boldsymbol{\beta}_\rho) \leq \right. \\ &\quad \left. \frac{\text{qchisq}(\alpha, k)}{\sum_{i=1}^n \rho^{(2)} \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}_\rho}{\sigma_n} \right)} \frac{\sum_{i=1}^n \rho^{(1)^2} \left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}}{\sigma_n} \right)}{\left(\frac{y_i - \mathbf{x}_i^t \boldsymbol{\beta}_\rho}{\sigma_n} \right)} \right\}. \end{aligned} \quad (48)$$

3.4 Least squares regression

The method goes through for the least squares functional with the advantage that explicit expressions are available. The result corresponding to (47) is

$$\mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_2) = \left\{ \boldsymbol{\beta} : \|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}\|_2^2 - \|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}_{2,n}\|_2^2 \leq \frac{\|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}\|_2^2}{n} \text{qchisq}(\alpha, k) \right\} \quad (49)$$

which is the same as

$$\mathcal{NS}(\mathbf{y}_n, \mathbf{x}_n, \alpha, T_2) = \left\{ \boldsymbol{\beta} : (\boldsymbol{\beta} - \boldsymbol{\beta}_{2,n})^t \mathbf{x}_n^t \mathbf{x}_n (\boldsymbol{\beta} - \boldsymbol{\beta}_{2,n}) \leq \frac{\|\mathbf{y}_n - \mathbf{x}_n \boldsymbol{\beta}_{2,n}\|_2^2 \text{qchisq}(\alpha, k)}{n - \text{qchisq}(\alpha, k)} \right\}. \quad (50)$$

where $\boldsymbol{\beta}_{2,n} = T_2(\mathbb{P}_n)$. The region is asymptotically equivalent to a standard α -confidence region for the ‘true’ parameter value.

3.5 Covering properties

The concept of a non-significance region makes no mention of a model or true values. Nevertheless there are situations where a model and its parameters are well founded and relate to well-defined properties of the real world. In such cases there is an interest in specifying a region which includes the real world value with the required frequency in repeated measurements. It has to be kept in mind however that covering true parameter values in simulations is not the same as covering the corresponding real values for real data (see Chapter 5.5 of Davies (2014), Stigler (1977), Chapter 8.1 of Hampel et al. (1986), Kunsch et al. (1993)).

Given this there is an interest in the covering properties of non-significance regions. Table 7 gives the frequencies with which the non-significance intervals (36) and the confidence intervals based on the rank statistics cover the population median and also the average lengths of the intervals. The results are for the normal, Cauchy, χ_1^2 and the Poisson $\mathfrak{Po}(4)$ distributions and four different sample sizes $n = 10, 20, 50, 100$ and are based on 1000 simulations. The discreteness of Poisson distribution was taken into account in the calculations of the non-significance region as follows. If a non-significance interval $[\ell, u]$ contains

	n	10	20	50	100
$N(0, 1)$	(36)	0.940 1.512	0.954 1.040	0.948 0.648	0.942 0.464
	rank	0.968 2.046	0.968 1.198	0.970 0.767	0.964 0.530
$C(0, 1)$	(36)	0.960 3.318	0.956 1.670	0.960 0.958	0.952 0.629
	rank	0.978 5.791	0.950 1.850	0.968 1.069	0.964 0.700
χ_1^2	(36)	0.944 1.368	0.936 0.877	0.932 0.550	0.942 0.396
	rank	0.982 2.064	0.958 1.086	0.970 0.675	0.968 0.452
$\text{Pois}(4)$	(36)	0.934 1.918	0.925 0.993	0.926 0.288	0.938 0.071
	rank	0.996 3.948	0.964 2.342	0.997 1.573	1.000 1.085

Table 7: Covering frequencies and average interval lengths based on 1000 simulations for the median for the 0.95-non-significance intervals as defined by (36) and (37) with $Z = N(0, 1)$ and the 0.95-confidence intervals based on the ranks. For each sample size the first column gives the covering frequency and the second the average interval length.

an integer it is by $[[\ell], \lfloor u \rfloor]$. If it does not contain an integer it is replaced by $[[\ell], \lceil u \rceil]$. The covering frequencies and lengths refer to this modified interval. In this well defined situation Table 7 indicates that the 0.95-non-significance intervals also have covering probabilities of about 0.95. The finite sample behaviour seems to be better than that of the ranks procedure. Both methods have approximately the correct covering frequencies but the lengths of the non-significance intervals are uniformly smaller than the lengths of the confidence intervals.

There is some theoretical explanation as to why the non-significance regions have covering frequencies given by α , at least asymptotically. Consider firstly i.i.d. integer valued random variables Y_j with a unique median ν . Then for a large sample size n

$$\sum_{j=1}^n |Y_j - \nu - bZ_j|$$

is, with large probability, minimized by putting $b = 0$. In other words the 0.95-non-significance interval is simply $[\nu, \nu]$ with a covering probability tending to one. This is illustrated by the Poisson distribution in Table 7.

Suppose that the Y_j are continuous random variables with median 0 and a density f which is continuous at 0 with $f(0) > 0$. Then the approximation

$$\sum_{i=1}^n \left| Y_i - \frac{bZ_i}{\sqrt{n}} \right| \approx \sum_{i=1}^n |Y_i| - bN(0, 1) + f(0)b^2 \quad (51)$$

holds (see the Appendix for a heuristic proof) and minimizing over b gives

$$\inf_b \sum_{i=1}^n \left| Y_i - \frac{bZ_i}{\sqrt{n}} \right| \approx \sum_{i=1}^n |Y_i| - \frac{\chi_1^2}{4f(0)}. \quad (52)$$

Moreover the same proof gives

$$\inf_b \sum_{i=1}^n \left| Y_i - \text{med}(\mathbf{Y}_n) - \frac{\theta}{\sqrt{n}} - \frac{bZ_i}{\sqrt{n}} \right| \approx \sum_{i=1}^n |Y_i - \text{med}(\mathbf{Y}_n)| + f(0)\theta^2 - \frac{\chi_1^2}{4f(0)} \quad (53)$$

from which the asymptotic α -non-significance interval

$$\left[\text{med}(\mathbf{Y}_n) - \sqrt{\frac{\text{qchisq}(\alpha, 1)}{4f(0)^2n}}, \text{med}(\mathbf{Y}_n) + \sqrt{\frac{\text{qchisq}(\alpha, 1)}{4f(0)^2n}} \right] \quad (54)$$

as defined in (36) and (37) follows. This latter interval is the same as the asymptotic confidence interval based on the median. Just as for the inverse rank method it does not require an estimate of $f(0)$.

L_1 linear regression can be treated in the same manner. Corresponding to (51) one has

$$\sum_{i=1}^n \left| Y_i - \frac{\mathbf{Z}_i^t \mathbf{b}}{\sqrt{n}} \right| \approx \sum_{i=1}^n |Y_i| - N(0, \mathbf{I}_k)^t \mathbf{b} + f(0)\|\mathbf{b}\|_2^2. \quad (55)$$

Applying this to the L_1 regression functional gives

$$\inf_{\mathbf{b}} \sum_{i=1}^n \left| Y_i - \mathbf{x}_i^t \boldsymbol{\beta}_{1,n} - \frac{\mathbf{x}_i^t \boldsymbol{\theta}}{\sqrt{n}} - \frac{\mathbf{Z}_i^t \mathbf{b}}{\sqrt{n}} \right| \approx \sum_{i=1}^n |Y_i - \mathbf{x}_i^t \boldsymbol{\beta}_{1,n}| + f(0)\boldsymbol{\theta}^t \mathbf{Q}_n \boldsymbol{\theta} - \frac{\chi_k^2}{4f(0)^2} \quad (56)$$

where $\mathbf{Q}_n = \frac{1}{n} \mathbf{x}_n^t \mathbf{x}_n$. From this the asymptotic α -non-significance region

$$(\boldsymbol{\beta} - \boldsymbol{\beta}_{1,n})^t \mathbf{Q}_n (\boldsymbol{\beta} - \boldsymbol{\beta}_{1,n}) \leq \frac{\text{qchisq}(\alpha, k)}{4f(0)^2n} \quad (57)$$

follows. It is the same as the α -confidence region based on the L_1 regression estimate $\boldsymbol{\beta}_1$, see for example Zhou and Portnoy (1996).

Table 8 gives the covering frequencies and average interval lengths for data generated according to

$$Y = -39.69 + 0.832 \cdot \text{Air.Flow} + 0.574 \cdot \text{Water.Temp} - 0.061 \cdot \text{Acid.Conc} + \varepsilon \quad (58)$$

using the L_1 coefficients for the stack loss data. The sample size is $n = 21$. The following four distributions for the error term ε are used: $\varepsilon = N(0, 1) * \text{Res}$, $\varepsilon = \sigma N(0, 1)$ for the normal distribution, $\varepsilon = \sigma L^*$ for the Laplace distribution and $\varepsilon = \sigma C^*$ for the Cauchy distribution where L^* and C^* are respectively the Laplace and Cauchy distributions closest to the $N(0, 1)$ distribution, Res are the residuals and σ the mean absolute deviation of the residuals of the stack loss data.

		β_2	β_3	β_4
residuals	(45)	0.944 0.265	0.982 0.682	0.998 0.248
	rank	0.976 0.390	0.970 1.205	0.970 0.273
Normal	(45)	0.954 0.381	0.946 1.042	0.964 0.442
	rank	0.974 0.435	0.956 1.208	0.962 0.542
Laplace	(45)	0.953 0.501	0.959 1.375	0.952 0.580
	rank	0.966 0.594	0.959 1.697	0.960 0.761
Cauchy	(45)	0.928 1.467	0.942 4.052	0.936 1.731
	rank	0.936 1.948	0.946 5.676	0.942 2.984

Table 8: Covering frequencies and average interval lengths for data generated according to (58) with different distributions for the error term: $\alpha = 0.95$.

Finally, in the case of non-linear L_1 regression the asymptotic α -non-significance is, under suitable regularity conditions, given by

$$(\beta - \beta_{nlr1,n})^t \mathbf{Q}_n (\beta - \beta_{nlr1,n}) \leq \frac{\text{qchisq}(\alpha, k)}{4f(0)^2 n} \quad (59)$$

where

$$\mathbf{Q}_n = \frac{1}{n} \sum_{i=1}^n \nabla_i \nabla_i^t$$

and

$$\nabla_i = \left(\frac{\partial m(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \theta_1}, \dots, \frac{\partial m(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \theta_k} \right)^t.$$

4 Choice of noise

It is possible to use random variables other than Gaussian. As an example the 0.95-non-significance intervals for the median of the stack loss data us-

ing $N(0, 1)$, ± 1 , $U(-1, 1)$, $\pm \text{beta}(5, 5)$ and the standard Cauchy distribution are (11.88, 18.63), (11.86, 18.25), (11.83, 18.16), (11.93, 18.21), (11.83, 18.16) and (11.00, 18.56) respectively. It is clear that the results depend on the choice of noise to some extent but that at least in this example the dependence is weak. Given the advantages of Gaussian noise are the easily available asymptotic expressions such (29) it would seem to be the default choice of noise.

Other possibilities are to make the noise dependent on the size of the covariates as in $W_{ij} = x_{ij}Z_{ij}$ or to randomly permute the covariates (see Anderson and Robinson (2001), Klingbiel (2009)).

5 Appendix

Consider

$$\sum_{i=1}^n \left| \varepsilon_i - \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right| \quad (60)$$

where the ε_i are symmetric, i.i.d. random variables with a continuously differentiable density at $u = 0$ with $f(0) > 0$. The $\mathbf{U}_{i\cdot}$ are k dimensional random variables $\mathbf{U}_{i\cdot} = (U_{i,1}, \dots, U_{i,k})^t$ where the U_{ij} are symmetric i.i.d. random variable with unit variance. The sum (60) may be decomposed as

$$\begin{aligned} \sum_{i=1}^n \left| \varepsilon_i - \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right| &= \sum_{\varepsilon_i \leq -\left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right|} (-\varepsilon_i + \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}}) \\ &\quad + \sum_{\varepsilon_i \geq \left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right|} (\varepsilon_i - \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}}) \\ &\quad + \sum_{|\varepsilon_i| \leq \left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right|} \left| \varepsilon_i + \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right| \\ &= \sum_{i=1}^n |\varepsilon_i| + \sum_{i=1}^n \pm \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} - \sum_{|\varepsilon_i| \leq \left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right|} \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \\ &\quad + \sum_{|\varepsilon_i| \leq \left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right|} \left(\left| \varepsilon_i - \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right| - |\varepsilon_i| \right). \end{aligned}$$

The random variables

$$V_i = \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \left\{ |\varepsilon_i| \leq \left| \frac{\mathbf{U}_{i\cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right\}$$

are i.i.d with mean zero and variance

$$\frac{1}{n} \mathbf{E}_U \left((U_{i \cdot}^t \mathbf{b})^2 \left(F \left(\left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right) - F \left(- \left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right) \right) \right) = o \left(\frac{\|\mathbf{b}\|_2^2}{n} \right).$$

This together with the central limit theorem implies

$$\sum_{i=1}^n \left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| = \sum_{i=1}^n |\varepsilon_i| + \mathbf{Z}^t \mathbf{b} + \sum_{|\varepsilon_i| \leq \left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right|} \left(\left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| - |\varepsilon_i| \right) + o(\|\mathbf{b}\|_2^2)$$

where $\mathbf{Z} \stackrel{D}{=} N(0, I_k)$. Denote the distribution function of $U_{i \cdot}^t \mathbf{b}$ by H . Then

$$\mathbf{E}_U \left(\left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \left\{ |\varepsilon_i| \leq \left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right\} \right) = \frac{2}{\sqrt{n}} \int_0^\infty w \left\{ |\varepsilon_i| \leq \frac{w}{\sqrt{n}} \right\} dH(w)$$

and taking the expected value with respect to ε_i gives

$$\begin{aligned} \mathbf{E} \left(\left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \left\{ |\varepsilon_i| \leq \left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right\} \right) \\ = \frac{2}{\sqrt{n}} \int_0^\infty w \left(F \left(\frac{w}{\sqrt{n}} \right) - F \left(-\frac{w}{\sqrt{n}} \right) \right) dH(w) \\ \approx \frac{4f(0)}{n} \int_0^\infty w^2 dH(w) = \frac{2f(0)\|\mathbf{b}\|_2^2}{n} \end{aligned}$$

as the U_{ij} are symmetric random variables with variance 1. A similar calculation gives

$$\mathbf{E} \left(|\varepsilon_i| \left\{ |\varepsilon_i| \leq \left| \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \right\} \right) \approx \frac{f(0)\|\mathbf{b}\|_2^2}{n}.$$

Putting this together leads to

$$\sum_{i=1}^n \left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \approx \sum_{i=1}^n |\varepsilon_i| + \mathbf{Z}^t \mathbf{b} + \frac{f(0)\|\mathbf{b}\|_2^2}{n}$$

and minimizing over \mathbf{b} results in

$$\inf_{\mathbf{b}} \sum_{i=1}^n \left| \varepsilon_i - \frac{U_{i \cdot}^t \mathbf{b}}{\sqrt{n}} \right| \approx \sum_{i=1}^n |\varepsilon_i| - \frac{\chi_k^2}{4f(0)}$$

where χ_k^2 is a chi-squared random variable with k degrees of freedom.

References

Akaike, H. (1973), Information theory and an extension of the maximum likelihood principle, in B. Petrov and F. Csaki, eds, ‘Second international symposium on information theory’, Akademiai Kiado, Budapest, pp. 267–281.

- Akaike, H. (1974), ‘A new look at the statistical model identification’, *IEEE Transactions on Automatic Control* **19**, 716–723.
- Akaike, H. (1981), ‘Likelihood of a model and information criteria’, *Journal of Econometrics* **16**, 3–14.
- Anderson, M. and Robinson, J. (2001), ‘Permutation tests for linear models’, *Australian and New Zealand Journal of Statistics* **43**(1), 75–88.
- Berk, R., Brown, L., Buja, A., Zhang, K. and Zhao, L. (2013), ‘Valid post-selection inference’, *Annals of Statistics* **41**(2), 802–837.
- Brownlee, K. A. (1960), *Statistical Theory and Methodology in Science and Engineering*, 2nd edn, Wiley, New York.
- Claeskens, G. and Hjort, N. L. (2003), ‘Focused information criterion’, *Journal of the American Statistical Association* **98**, 900–916.
- Davies, L. (2014), *Data Analysis and Approximate Models*, Monographs on Statistics and Applied Probability 133, CRC Press.
- Dutter, R. (1977a), ‘Algorithms for the Huber estimator in multiple regression’, *Computing* **18**, 167–176.
- Dutter, R. (1977b), ‘Numerical solution of robust regression problems: Computational aspects, a comparison’, *Journal of Statistical Computing and Simulation* **5**, 207–238.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J. and Stahel, W. A. (1986), *Robust Statistics: The Approach Based on Influence Functions*, Wiley, New York.
- Hosmer, D. W. and Lemeshow, S. (1989), *Applied Logistic Regression*, Wiley, New York.
- Huber, P. J. and Ronchetti, E. M. (2009), *Robust Statistics*, second edn, Wiley, New Jersey.
- Klingbiel, D. (2009), Permutation Tests Applied to Multivariate Regression, PhD thesis, Philosophisch-naturwissenschaftliche Fakultät der Universität Bern, Bern, Switzerland.

- Koenker, R. (2010), ‘quantreg: Quantile regression’, <http://CRAN.R-project.org/package=quantreg>. R package version 4.53.
- Koenker, R. and Portnoy, S. (1997), ‘The Gaussian hare and the Laplacian tortoise: computability of squared-error versus absolute-error estimators’, *Statist. Sci.* **12**(4), 279–300.
- Kunsch, H., Beran, J. and Hampel, F. (1993), ‘Contrasts under long-range correlations’, *Annals of Statistics* **21**, 943–964.
- Lindsay, B. and Liu, J. (2009), ‘Model assessment tools for a model false world’, *Statistical Science* **24**(3), 303–318.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (2003), *Numerical Recipes in Fortran 77: The Art of Scientific Computing*, Vol. 1, second edn, Cambridge University Press.
- R Core Team (2013), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: <http://www.R-project.org/>
- Schwarz, G. E. (1978), ‘Estimating the dimension of a model’, *Annals of Statistics* **6**(2), 461–464.
- Stigler, S. M. (1977), ‘Do robust estimators work with real data? (with discussion)’, *Annals of Statistics* **5**(6), 1055–1098.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *J. Royal. Statist. Soc. B.* **58**(1), 267–288.
- Zhou, K. Q. and Portnoy, S. L. (1996), ‘Direct use of regression quantiles to construct confidence sets in linear models’, *Annals of Statistics* **24**(1), 287–306.